

AFRL-IF-RS-TR-2001-91
Final Technical Report
May 2001



FUSION OF STOCHASTIC AND LINGUISTIC INFORMATION USING A CONDITIONAL EVENT FRAMEWORK

University of Massachusetts

Patrick A. Kelly, Haluk Derin and Wei-Bo Gong


APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

20010713 045

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2001-91 has been reviewed and is approved for publication.



APPROVED: MARK G. ALFORD
Project Engineer



FOR THE DIRECTOR:

JOSEPH CAMERA, Chief
Information & Intelligence Exploitation Division
Information Directorate

If your address has changed or if you wish to be removed from the Air Force Research Laboratory Rome Research Site mailing list, or if the addressee is no longer employed by your organization, please notify AFRL/IFEA, 32 Brooks Road, Rome, NY 13441-4114. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE May 01		3. REPORT TYPE AND DATES COVERED Final Sep 98 - Jun 00
4. TITLE AND SUBTITLE FUSION OF STOCHASTIC AND LINGUISTIC INFORMATION USING A CONDITIONAL EVENT FRAMEWORK			5. FUNDING NUMBERS C - F30602-98-C-0263 PE - 62702F PR - 4600 TA - II WU - D2	
6. AUTHOR(S) Patrick A. Kelly, Haluk Derin and Wei-Bo Gong				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Massachusetts College of Engineering Box 35110 Amherst, MA 01003-5110			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL/IFEA 32 Brooks Road Rome, NY 13441-4114			10. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2001-91	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Mark A. Alford, IFEA, 315-330-3802				
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report addresses a framework for merging different types and sources of information for applications involving image estimation. Complex decision-making systems rely on feature estimates derived from image data. It is imperative that all available information be used to effectively generate high quality estimates. This information includes stochastic raw sensor data, conditional information obtained from other sources and linguistic information such as if-then rules obtained from human experts who supervise the processing.				
14. SUBJECT TERMS conditional event algebra, random sets, bayesian networks, image processing, linguistic processing			15. NUMBER OF PAGES 28	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

TABLE OF CONTENTS

I. EXECUTIVE SUMMARY	1
II. OBJECTIVE AND SCOPE OF WORK	3
III. PROJECT TASK SUMMARIES	4
1. Development of a conditional event framework for fusion of stochastic and linguistic information	4
2. Development of algorithms for iterated image estimation and image model updating	5
3. Development of algorithms for knowledge-aided model identification and optimization	10
IV. CONCLUSIONS	14
V. REFERENCES	15

LIST OF FIGURES

Figure 1 Edge Detection	7
Figure 2 Boundary Localization	7
Figure 3 Active Contour Matching	7
Figure 4 (a) Network for multi-target scene; (b) Probability propagation	9
Figure 5 (a) True targets; (b) Observed image.	10
Figure 6 Targets estimated by network	10
Figure 7 Adaptive system model	13
Figure 8 Adaptive algorithm applied to restoration of text image	14

LIST OF ACRONYMS

PS-CEA	PRODUCT SPACE CONDITIONAL EVENT ALGEBRA
RS	RANDOM SET
GUI	GRAPHICAL USER INTERFACE
LV	LEFT VENTRICLE
2DE	TWO DIMENSIONAL ELECTROCARDIOGRAPHIC
MAP	MAXIMUM-A-POSTERIORI
LMS	LEAST MEAN SQUARE

I. EXECUTIVE SUMMARY

This is the final project report for Contract F30602-98-C-0263, titled "Fusion of Stochastic and Linguistic Information Using a Conditional Event Framework". The project addressed a fundamental and challenging problem in information fusion, namely, the development of a framework for merging different types and sources of information for applications involving image estimation. The complex decision-making systems characteristic of modern military, medical and industrial applications frequently rely on feature estimates derived from image data. In such demanding applications, it is imperative that all available information be used effectively to generate high-quality estimates. In typical systems, this information can include the *stochastic* raw sensor data, *conditional* information such as preliminary estimates obtained from other sources, and *linguistic* information such as if-then rules supplied by human experts who supervise the processing. Without a systematic and consistent framework for dealing with these diverse types of information, it is very difficult to design near-optimal fusion-based algorithms.

In this project, our aim was to develop an information fusion framework by making use of recent advances in the theory of *Product Space Conditional Event Algebras* (PS-CEA's) and *Random Sets* (RS's) [1,2,3,4,10]. PS-CEA and RS techniques hold the promise of enabling stochastic, conditional and linguistic information to be represented in the common setting of a single probability space. In principle, this allows the computation of one joint distribution encompassing the diverse information sources, and thus allows standard statistical techniques such as Bayesian estimation algorithms to be generalized and applied to fusion-based tasks. As a proof-of-concept for this approach to information fusion, we considered its application to three concrete tasks:

1. Iterated image estimation: Development of algorithms for estimate updating based on sequences of data from imaging sensors.
2. Image model updating with uncertain observations: Development of algorithms that use uncertain information about a scene to update conditional models.
3. Knowledge-aided model identification and optimization: Use of expert-supplied rules together with image data to identify and optimize processing systems.

Results of our work on these tasks are described below, and in more detail in the published papers and thesis. Based on these results, we are able to state the following general conclusions:

1. The PS-CEA approach is helpful for representing some types of partial statistical knowledge. For example, one common problem is that we know an observation

distribution only when it is conditioned on the validity of some previous estimate. If we also have confidence values assigned to estimate outcomes, then the PS-CEA representation can be used to formulate algorithms for iterated estimation. These can have applications to, for example, scene estimation from a sequence of observed images. In general, however, the PS-CEA representation appears to be of somewhat limited use for the applications that we considered. One problem is that, while in principle one can write a joint distribution for any number of product space conditional events, in practice the computation of that distribution is very difficult when the number of conditional events exceeds two. Another problem is that the very common image estimation operation of conditioning on a continuous-valued observation (i.e., conditioning on an event whose probability equals or approaches zero) does not appear to be handled well by a PS-CEA representation. Hence, although PS-CEA can provide a conceptual framework for information fusion, for image estimation applications it seems that it is necessary to consider alternate representations for some types of information and computations. (See the discussion in Sections III.1 and III.2 below.)

2. One of the reasons for considering operations like iterated estimation for image processing is that nominally optimal Bayesian estimates are too computationally complex to implement directly. Similar difficulties are encountered in decoding problems in communication systems, and for some of those problems it has recently been found that effective and practical iterated computations can be organized via the framework of *probability propagation* in *Bayesian networks* [5]. In those networks, conditioning on observations is represented by impulse functions. We did a preliminary study of this approach for organizing the computations in image estimation, and believe that it is potentially useful for applications such as target detection and recognition. (See Section III.2.)
3. For computations involving multiple conditional events such as if-then rules, the RS representation for rules proposed by Mahler [1,2,3,4] does lead to a practical and useful computational framework for fusion with data. We considered application of the framework to a problem of user-assisted image processor optimization, and believe that it can be generally applicable in problems involving expert-guided processing. (See Section III.3.)

The rest of this report is organized as follows: For reference, the Objective, Scope, and Tasks/Technical Requirements of the project are first restated from the Contractor Statement of Work in Section II. Section III contains summary descriptions of the approaches, results and conclusions of the three main project tasks. Section IV contains concluding remarks, followed by the bibliography in Section V. This report resulted in several three conference papers that describe task outcomes in detail [6,7,8]. Other results include a Masters thesis developed from work on the knowledge-aided identification and

optimization task [9] and software for a graphical user interface (GUI) implementation of the algorithm developed in the thesis.

II. OBJECTIVE AND SCOPE OF WORK

For reference, the Objective, Scope, and Tasks/Technical Requirements of the project are restated below from the Contractor Statement of Work.

1.0 OBJECTIVE:

1.1 The objective of the proposed effort is to investigate and develop algorithms and prototype software demonstrations establishing the feasibility and utility of conditional event algebra techniques for information fusion, with specific applications to (a) iterated image estimation; (b) image model updating with uncertain observations; and (c) knowledge-aided model identification and optimization.

2.0 SCOPE:

2.1 The scope of this effort is to investigate new approaches for information fusion and to demonstrate their usefulness in image data processing for decision-making systems.

4.0 TASKS/TECHNICAL REQUIREMENTS:

4.1 The contractor shall accomplish the following:

4.1.1 Develop a conditional event framework for fusion of stochastic and linguistic information.

4.1.2 Develop algorithms for iterated image estimation and image model updating, using the framework developed under 4.1.1.

4.1.3 Develop algorithms for knowledge-aided model identification and optimization, using the framework developed under 4.1.1.

4.1.4 Identify appropriate test systems for algorithms developed under 4.1.2 and 4.1.3.

4.1.5 Test and validate the algorithms developed under 4.1.2 and 4.1.3.

4.1.6 Provide a final report documenting all work accomplished under the project.

4.1.7 Continually provide status reports as requested by the Program Manager of this effort.

4.1.8 Conduct presentations/meetings at times and places specified in the Schedule.

III. PROJECT TASK SUMMARIES

1. Development of a Conditional Event Framework for Fusion of Stochastic and Linguistic Information

Under this task, we developed a Generalized Bayesian approach for information fusion that is based on the PS-CEA iterated conditioning formula described in [1,10]. The fundamental equation for the Generalized Bayesian approach is the following:

$$P((X|A)|(Y|B)) = \frac{P(X, A, B^c) + P(B)\{P(X, A | Y, B) + P(A^c | Y, B)P(X | A)\}}{P(A \cup B)} \quad (1)$$

In (1), X is the quantity to be estimated. The conditioning event A represents previously available information (e.g., operating rules or prior estimates). The conditional event $(X|A)$ then represents the "prior" model (i.e., the model that encompasses the available information before the arrival of a new observation). The event $(Y|B)$ represents the new observation, also possibly conditioned on some other event. The operation implemented by (1) is conditioning of the prior information on newly-arriving information to generate an updated (posterior) model used for estimation. The Generalized Bayesian approach results in recursive, iterated-conditioning estimators that (in a manner similar to Kalman Filtering) use the posterior distribution at one stage to generate the conditioning event for the prior at the next stage. The key considerations in using the approach in specific applications are the definitions of the conditioning events and the computation of the component distributions on the right-hand side of (1). Throughout this project we examined estimators based on (1), or when the computations in (1) proved infeasible, on alternatives in the same spirit of iterated conditioning.

As one potential application we considered the general problem of image sequence estimation, in which the information passed from one stage to the next is the current estimate of the image and/or its distribution, together with a measure of confidence in the estimate accuracy. When the next observed frame arrives, the problem is to fuse the new observation and the previous estimate to obtain an updated estimate for the current frame. Using (1), we derived a distribution update formula that could be used, for example, to generate improved scene estimates from sequentially-arriving observations. The update formula has the form:

$$\begin{aligned}
& P((X^{(k+1)} | (\hat{X}^{(k)} | C^{(k)})) | y^{(k+1)}) \\
& = P(X^{(k+1)} | \hat{X}^{(k)}, y^{(k+1)})P(C^{(k)}) + P(X^{(k+1)} | y^{(k+1)})(1 - P(C^{(k)}))
\end{aligned} \tag{2}$$


where $X^{(k+1)}$ is the current scene, $\hat{X}^{(k)}$ is the scene estimate from the previous frame, $C^{(k)}$ is an event under which the estimate $\hat{X}^{(k)}$ is accurate, $P(C^{(k)})$ is the confidence in that estimate, and $y^{(k+1)}$ is the new (current) observation. (A similar formula can be derived when there is distribution information passed from the previous frame, instead of just an estimate value.) Equation (2) is a fusion formula, generating the updated posterior scene distribution given the current observation and the conditional (partial) information in the previous estimate (conditioned on its accuracy). As (2) shows, the updated distribution is actually a weighted combination of two posterior distributions: one given the observed data and that the previous estimate is accurate, and the second given the observed data alone. The development of the Generalized Bayesian approach, its possible applications in image estimate updating, and some tests on simulated image sequences are described in more detail in the conference paper [6].

2. Development of Algorithms for Iterated Image Estimation and Image Model Updating

Having derived a conceptual Generalized Bayesian approach and considered its possible application to image sequence estimation, we next examined how well it might actually apply to two concrete problems of iterated image estimation and image model updating. Problems of iterated image estimation arise whenever we have a time-sequence of observed images that is to be used in estimating a stationary or moving scene. In [3], we described how the Generalized Bayesian formula (1) could be applied to estimating a static scene from a sequence. In this task, we wanted to extend the approach to the estimation of a moving scene. The test application that we chose for this was locating the moving boundary of the left ventricle (LV) of the human heart in a sequence two-dimensional echocardiographic (2DE) images. Our chief reason for choosing this as a test case is that the ultrasound image quality in any given frame can be very poor, so that accurate estimation from current data alone is difficult. We need to fuse the data with other available information such as expert knowledge of LV shape or a prior good-quality estimate in order to generate sufficiently good estimates from the inevitable poor-quality frames. Also, the ability to accurately estimate and track the LV boundary would have practical significance in clinical applications (see, e.g., [11,12]), and data sequences are readily available for testing. In order to use the image estimate formula (2) that we derived from the Generalized Bayesian approach, we need to compute $P(X^{(k+1)} | \hat{X}^{(k)}, y^{(k+1)})$; that is, we need to be able to find the distribution of the current scene, given the new observation and that the previous estimate is accurate. We formulated an approach to this based on *deformable template* models [13] and *snakes*

[14]. In that approach, the previous estimate and the observation are weighted in an energy function that determines the conditional distribution for the current scene. If $P(C^{(k)})=1$, then by maximizing that conditional distribution we obtain the desired current-frame estimate. An example application of this approach is shown in Figures 1-3. From left to right, Figure 1 shows a typical image frame, the edges detected in that frame, and the edges superimposed on the observation. This illustrates some of the difficulties in processing 2DE images – because of poor image contrast and tissue variations in the scene, there are numerous missing and false edge pixels. Our objective is to find the best closed-boundary estimate for the LV (the dark region in the middle of the frame). Figure 2 shows a prototype template (derived from a high-quality estimate from a previous frame) on the left and the results of a localization procedure that modifies this template and matches it to the observed image edges on the right. Finally, Figure 3 shows how the boundary estimate is refined using snake techniques, with the final closed boundary estimate shown superimposed on the observed frame on the right. This illustrates that even relatively inaccurate templates derived from temporally distant frames can provide significantly helpful information for overcoming a poor-quality observation to produce a good LV boundary estimate.

The outcome shown in Fig. 3(b) is the result of maximizing $P(X^{(k+1)} | \hat{X}^{(k)}, y^{(k+1)})$ with our model. This already performs information fusion, since it merges the prior knowledge of how the LV boundary is expected to move from frame to frame with the actual observed image. To fully use the formula (2) that follows from the Generalized Bayesian approach, we must define a confidence level in the accuracy of the previous-frame LV boundary estimate. When this is done properly, then the boundary matching procedure will be improved by properly weighting the contributions of the previous estimate and the data. That is, through the second term on the right-hand side of (2), the current estimate will be pulled away from a value determined solely by the first term (which is what is shown in Fig. 3) to a greater or lesser extent depending on the confidence. In our ongoing work on this task we are exploring how best to define and use confidence measures. The results of the work on this task will be described more fully in a paper to be submitted to the SPIE Conference on Medical Imaging in February, 2001.



(a) Input frame

(b) Edge map

(c) Superimposed edges

Figure 1 Edge Detection




(a) Prototype template

(b) Edge map of input frame

(c) Localization

Figure 2 Boundary Localization



(a) Placement of snake

(b) Estimate of LV boundary

(c) Superimposed result

Figure 3 Active Contour Matching

The second concrete application that we investigated under this task is image model updating for target detection and recognition. To perform effectively, these algorithms must use both prior knowledge and observations. In particular, there has been much recent research on defining or approximating *maximum a posteriori* (MAP) estimates for detection and recognition (e.g., [18,19]). These estimates are based on maximizing posterior distributions which are typically very complex. Feasible algorithms usually approximate this maximization through iterative computations of partial estimates (e.g., using annealing approaches which update part of the estimate given that the rest of the current estimate is accurate). It can be difficult to determine exactly what sequence of partial estimate computation will be most practical and effective. Also, while these algorithms generate estimates, for fusion with other information it would actually be more useful to have probabilities or confidence measures for a set of plausible alternatives. Our goal under this task was to investigate the possible application of the Generalized Bayesian approach to determine the best sequence of image model (i.e., distribution) updating for defining posterior target shape and label probabilities.

We immediately faced a difficulty in trying to apply (1) directly to this problem, however, which is that we need to represent a conditional event in which conditioning is on a continuous-valued observation. (This is in contrast to the iterated estimation problem described above, where the conditional event has conditioning on $C^{(k)}$, a binary-valued event representing validity of the previous estimate.) Formula (1) weighs probability terms involving both the conditioning event and its complement. The problem is how to interpret the event that we observe some point value of a continuous-valued random field. In particular, if we say that the probability that the field takes any point value is zero, then the complement event has probability one and only the term involving the complement appears in the computation, which does not lead to a useful result. This is illustrative of what appears to be a limitation of PS-CEA -- we were not able to find a satisfactory way to use PS-CEA to represent conditioning on continuous-valued observations.

In searching for alternative frameworks for representing model updating from partial information, we noted that problems of finding effective, computationally feasible approximations to posterior distributions and MAP estimates also arise in communication system applications. In particular, there has been a great deal of recent interest in iterative soft-decision decoders (e.g., *turbocoders* [15]) that are based on updating of symbol probabilities. It has been noted that the types and sequence of computations involved in iterative soft-decision decoding can be organized via the framework of *probability propagation* in *Bayesian networks* [5,15]. In a Bayesian network, nodes representing variables are connected by directed links representing conditioning. For singly-connected networks (i.e., networks having only one path (when link directions are ignored) between any two nodes), there is an iterative algorithm for propagating updated probabilities forward and backward along the network that can be used to generate exact

posterior probabilities at nodes. In multiply-connected networks, such as for turbocodes, application of the probability propagation algorithm does not produce exact posterior probabilities. However, as is noted in [5,15], the excellent decoding performance obtained by using probability propagation even for codes having multiply-connected networks suggests that it is an effective way for approximating MAP estimation.

In this task, we represented the variables and conditioning relations in a target detection and recognition problem using a (multiply-connected) Bayesian network, and considered the computations that result from applying a probability propagation algorithm to that network. Figure 4 shows the network that we used and the order of propagation for updated probabilities. In the figure, $p(S^{(k)})$ represents a prior distribution incorporating shape characteristics of a target of type k ; L is the label field (telling which target in the scene is of which type); and y is the observed image. The computation of updated target type and shape probability distributions based on this network is described in detail in the paper [7].




Figure 4 (a) Network for multi-target scene; (b) Probability propagation.

The test case considered in the paper used relatively simple target shape prior distributions and a very noisy observation. Even so, the algorithm was quite successful in estimating the target shapes, and it generated posterior target label probabilities that were very highly concentrated on the true labels. Figure 5 shows the true simulated targets on the left and the observed on the right. Figure 6 shows the recovered targets obtained by thresholding the posterior shape distributions for the most likely target labels found by the network. In our ongoing work, we are now looking at extending this approach to more informative and realistic target shape distributions, for example, edge-based models similar to those used in [18] or [19]. We are also investigating the use of this framework

for fusing multiple observations by having a large network for data fusion encompassing subnetworks of the form of Fig. 4 for the individual observations. We believe that there are two significant potential benefits from using the probability propagation approach. First, it gives us a practical framework for organizing computations to approximate MAP estimation; and second, the algorithm actually generates distributions on nodes which, even if not the precise posterior distributions, still give us more information for further processing (such as fusion with other data) than is provided by just an estimate value.




Figure 5 (a) True targets; (b) Observed image.




Figure 6 Targets estimated by network

3. Development of Algorithms for Knowledge-Aided Model Identification and Optimization

Modern decision-making systems for military, medical and industrial applications rely on information in image data. Often, the raw sensor data must be extensively processed in order to extract the most useful features. The exact form of the processing is determined by a set of parameters that are chosen to optimize some qualities of image features at the processor output. Image processors are often sufficiently complicated that it is difficult even for expert users to easily find the processing parameters that give the desired output for a given input image. (For example, see the discussion of the difficulties of manual processing selection for medical ultrasound systems in [11].) Furthermore, with the prevalence of image data in so many applications, image processors often must be applied by non-expert users who may have very little knowledge of how best to set the input parameters. One could avoid the difficulties of manual parameter selection by having an automated system that makes use of image and processor models together with a quantitative output performance measure (e.g., mean-squared error for restoration or probability of error for segmentation) to set parameters optimally. However, this approach has some significant drawbacks. For one, it is difficult to specify an image model that is sufficiently accurate that processing designed to be optimal for the model is close to optimal for a real observation. This has led to some interest in adaptive processors that adjust models according to measured input and output characteristics [16,17]. Since image processing is computationally expensive and time-consuming, any realistic adaptive system would have to start with a good enough model that the necessary adjustments could be made with relatively few iterations. Perhaps even more problematic is the specification of an appropriate performance measure. It is well-known that convenient quantitative measures such as mean-squared error fail to capture many visually important qualities, and that optimization with respect to such measures may not give an output that is most useful to a human observer [17]. It has been noted that what is really needed in many applications is a system for *qualitative optimization* of output features [17].

In fact, there is often a great deal of qualitative and linguistic information about processors and their effectiveness that is available. This information can take the form of operating rules describing how expert users adjust input parameters to improve the output quality, and user assessments of the visual qualities of the processed image. The key question is how to incorporate this qualitative information in a numerical optimization scheme for input parameter selection. In this task, we investigated the use of the RS representation for rules and other linguistic information described in [1,2,4] for qualitative optimization of an image processor. The test system that we considered was 2D isotropic Wiener filtering to restore noisy and blurred images. The reason for choosing this test case is that the input (filter) parameterization is relatively simple (only two parameters need to be specified), but how those parameters influence output visual quality is quite complex and non-obvious. We used three indicators of output quality: edge blur, edge ringing, and noisiness. For a given input image, we view the Wiener filter as a system that maps the input parameters to the output qualities. It is that system that we need to optimize. Since what is "optimal" will vary depending on the input and

the user's preferences, we cannot do this optimization with a fixed system model. Instead, we need an adaptive approach, recognizing that for practical considerations the adaptation can take only a few iterations, and that to match user preferences the adaptation needs to be driven by user-supplied quality assessments provided in some natural way.

A block diagram of the approach we used for this is shown in Figure 7. We first noted that it is difficult even for experts to state rules that directly relate filter input parameters to output image feature quality. However, it is relatively easy to state how *changes* in output features are generally related to changes in filter *properties* such as peak frequency, energy and maximum frequency-response slope. These filter properties are in turn related to filter input parameters through fixed numerical relations. So, in Fig. 7, System 1 is a fixed, nonlinear mapping from filter input parameters to filter properties. System 2 is an adaptive, user-determined linear mapping from changes in filter properties to changes in output image features. To have a good starting point, we used an initial System 2 model found from maximum-likelihood estimation applied to a distribution determined by random sets corresponding to seven expert-supplied operating rules. (For example, one of the rules was: "If filter energy is decreased and maximum frequency-response slope is decreased, then edge ringing will decrease and noisiness will decrease.") This nominal model is modified during processor operation to account for different input images and user preferences. When an input image is presented to the processor, the current system model is used to select the filter parameters. The resulting Wiener filter is applied, and the user assesses the output quality by indicating how each output feature needs to be improved (e.g., "less blur"). This linguistic information is converted to numerical values which are used to drive a *least mean-square* (LMS) adaptive filtering algorithm. This algorithm adjusts the System 2 model and uses the inverse of the adjusted model to find the desired changes in filter properties. These are added to the current filter property state to find the updated filter properties. Finally, the inverse of the System 1 model is applied to the new filter property state to select updated (and better) input parameters. The procedure continues until the user is satisfied with the output quality.

In tests on many different images, this approach has achieved a large degree of improvement in output quality after just a few iterations. An example is shown in Figure 8. The input image is shown in part (a). The output of a nominally-optimal Wiener Filter (the filter that would be applied in a non-adaptive implementation) is shown in part (b). The remaining figures show successive outcomes of the filter adjusted by the adaptive algorithm, with the process stopping after four iterations (with final output shown in part (f)). The development of the adaptive system approach is described more fully and with more examples in the paper [8] and the thesis [9]. The thesis also describes a GUI implementation of the adaptive system.

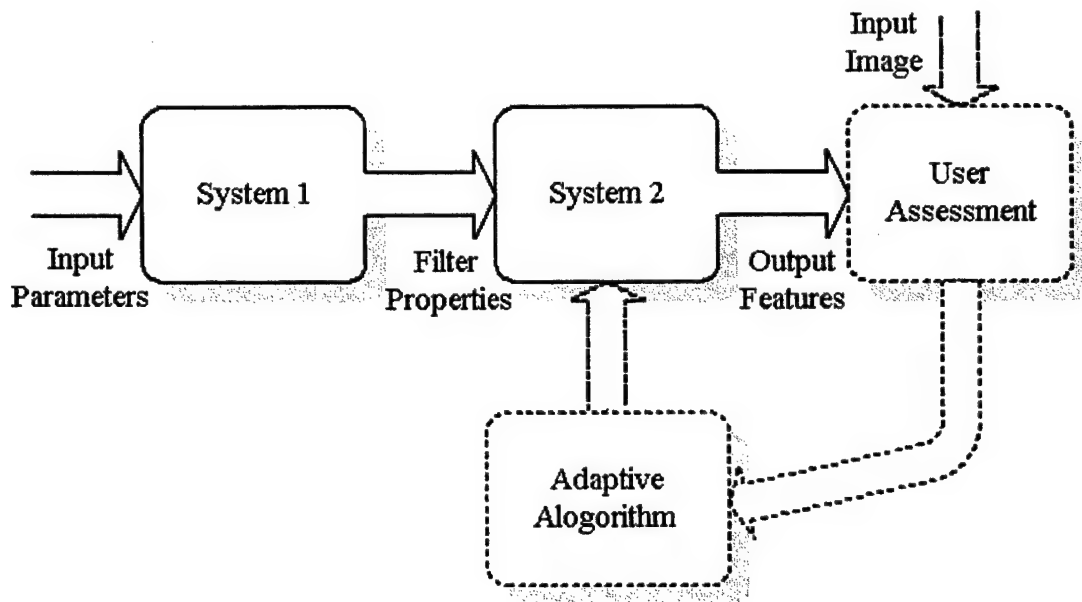


Figure 7 Adaptive system model.

While this approach generally succeeded in finding a good Wiener Filter, we do not view optimizing over this relatively simple class of filters as the most important aspect of the work. Rather, it is that we have a system that in the test case of Wiener filtering performs qualitative optimization, and that it does so with the only user intervention being the quality assessment - the filter parameter adjustment is done automatically, so that the user does not have to be an expert in the filter operation. We expect that with approaches similar to those developed under this task, systems can be developed for qualitative optimization of more complicated image processors.




Figure 8 Adaptive algorithm applied to restoration of text image.

IV. CONCLUSION

In this report, we have described the work done under Contract F30602-98-C-0263 for the period September 1998 – June 2000. As stated in Section II (from the Contractor Statement of Work), the scope of this effort was to investigate new approaches for information fusion and to demonstrate their usefulness in image data processing for decision-making systems. We have developed a framework based on Product Space Conditional Event Algebra (PS-CEA) representations for conditional events and applied it to iterated image estimation. We have investigated the use of an alternative framework based on probability propagation in Bayesian networks for organizing the computations in model updating for target detection and recognition. Random Set (RS) representations for rules and other linguistic information have been used to develop an approach to adaptive, qualitative optimization of image processor input parameters. The results that we have obtained are preliminary, and further work needs to be done under each of the tasks to produce practical working systems. However, even these preliminary results show the utility of incorporating available conditional or linguistic information in the data

processing. We believe that the results do demonstrate that it is worthwhile to continue to investigate approaches such as PS-CEA, RS, and Bayesian network techniques that allow diverse types of information to be included systematically and consistently in complex problems of image estimation.

V. BIBLIOGRAPHY

- [1] I. R. Goodman, R. P. S. Mahler and H. T. Nguyen, *Mathematics of Data Fusion*, Kluwer, 1997.
- [2] R. P. S. Mahler, "Representing Rules as Random Sets, I: Statistical Correlation Between Rules," *Information Sciences*, pp. 47-68, 1996.
- [3] R. P. S. Mahler, "Representing Rules as Random Sets, II: Iterated Rules," *International Journal of Intelligent Systems*, pp. 583-610, 1996.
- [4] R. P. S. Mahler, *An Introduction to Multisource-Multitarget Statistics and its Applications*, Lockheed-Martin Technical Report, 2000.
- [5] B. J. Frey, *Graphical Models for Machine learning and Digital Communication*, MIT Press, 1998.
- [6] P. Kelly, H. Derin and W. Gong, "Some Applications of Conditional Events and Random Sets for Image Estimation and System Modeling," *Proc. SPIE Conf. on Signal Processing, Sensor Fusion and Target Recognition VIII*, Orlando, April 1999.
- [7] P. A. Kelly, H. Derin and R. Zhao, "Target Detection and Recognition Using Markov Modeling and Probability Updating," *Proc. SPIE Conf. on Sensor Fusion: Architectures, Algorithms and Applications IV*, Orlando, April 2000.
- [8] P. A. Kelly, H. Derin and P. Vaidya, "Qualitative Optimization of Image Processing systems Using Random Set Modeling," *Proc. SPIE Conf. on Signal Processing, Sensor Fusion and Target Recognition IX*, Orlando, April 2000.
- [9] P. Vaidya, *Qualitative Optimization of Image Processing Systems*, M.S. thesis, Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, September 2000.
- [10] I. R. Goodman, "Toward a Comprehensive Theory of Linguistic and Probabilistic Evidence: Two New Approaches to Conditional Event Algebra," *IEEE Trans. Systems, Man & Cybernetics*, pp. 1685-1698, 1994.
- [11] I. Mikic, S. Krucinski and J. D. Thomas, "Segmentation and Tracking in Echocardiographic Sequences: Active Contours Guided by Optical Flow Estimates," *IEEE Trans. Medical Imaging*, pp. 274-284, 1998.
- [12] A. Fenster and D. B. Downey, "3-D Ultrasound Imaging: A Review," *IEEE Trans. On Engineering in Medicine and Biology*, pp. 41-51, 1996.
- [13] A. K. Jain, Y. Zhong, and S. Lakshmanan, "Object Matching Using Deformable Templates," *IEEE Trans. Pattern Analysis and Machine Intelligence*, pp. 267-278, 1996.

- [14] C. Xu and J. L. Prince, "Snakes, Shapes, and Gradient Vector Flow," *IEEE Trans. Image Processing*, pp. 359-369, 1998.
- [15] F. R. Kschischang and B. J. Frey, "Iterative Decoding of Compound Codes by Probability Propagation in Graphical Models," *IEEE Journal on Selected Areas in Communications*, pp. 219-230, 1998.
- [16] R. L. Lagendijk and J. Biemond, *Iterative Identification and Restoration of Images*, Kluwer, 1991.
- [17] M. R. Banham and A. K. Katsaggelos, "Digital Image Restoration," *IEEE Signal Processing Magazine*, pp. 24-41, 1997.
- [18] M. I. Miller, U. Grenander, J. A. O'Sullivan and D. L. Snyder, "Automatic Target Recognition Organized via Jump-Diffusion Algorithms," *IEEE Trans. Image Processing*, pp. 157-176, 1997.
- [19] S. Z. Der and R. Chellappa, "Probe-Based Automatic Target Recognition in Infrared Imagery," *IEEE Trans. Image Processing*, pp. 92-102, 1997.

**MISSION
OF
AFRL/INFORMATION DIRECTORATE (IF)**

*The advancement and application of Information Systems Science
and Technology to meet Air Force unique requirements for
Information Dominance and its transition to aerospace systems to
meet Air Force needs.*